OHEP Scientific Computing Report

2575 Sand Hill Road - MS 71 Menlo Park, CA 94025
Phone: 650-926-3824 Fax: E-Mail: richard@slac.stanford.edu

Introduction

Scientific computing at SLAC in the PPA and Accelerator (AD) Directorates currently support 3 active experiments (Fermi, EXO, ATLAS) and one experiment ramping down (BaBar). Other supported areas are HEP theory and cosmology simulations, and lepton collider detector studies. AD, of course, supports the advanced accelerator modeling. A Scientific Computing Applications department was formed in PPA in 2010 to leverage expertise in areas like data handling, GEANT4 and xrootd. Our research program will bring SCDMS online soon, and LSST is expected to become the flagship project. SLAC is getting experience via participation in DES and also pursuing opportunities in CTA.

Computing Facilities

Most of the hardware used by PPA is housed in the central computing center (CD) building, and is maintained by CD staff. Currently PPA experiments have some 8000 cores and 3 PB of disk. After considerable effort in building upgrades, there is now sufficient capacity in the CD building to house anticipated SLAC-wide growth in installed hardware until the joint SLAC-Stanford Research Computing Center opens circa 2013.

The Laboratory has adopted a recharge model for computing support which recognizes the importance of scientific computing by subsidizing costs to the tune of 20 cents on the dollar. The algorithm recovers cost for a fraction of the teams involved in supporting compute and storage clusters and is based on numbers of compute nodes and storage servers.

BaBar

The *BaBar* detector at the PEP II e^+e^- asymmetric collider at SLAC collected data at the Y(4S), Y(3S) and Y(2S) production thresholds from Oct 22nd 1999 to Apr 7th 2008.

BaBar, under the responsibility of the SLAC PPA Directorate, is now focusing on the full exploitation of the accumulated datasets. *BaBar* Computing is distributed among the host laboratory SLAC and four Tier A sites: CCIN2P3, Lyon, France (simulation production, skim production, data host), CNAF, Bologna, Italy (simulation production and data host), GRIDKA, Karlsruhe, Germany (skim production, data host), UVIC, Victoria, Canada (simulation production, data host). Simulation production is done also at ~15 other universities. INFN Padova continues to serve as the raw data backup.

After the end of data taking the data has undergone a major reprocessing using software releases of the 24 series. The production of R24 data is the final and most complete reconstruction of the recorded data. In 2009 many major skim releases have been produced and in 2010 an extra production of Monte Carlo samples of B meson pairs and charm mesons has brought these samples to the equivalent of 10 times the recorded luminosity. In 2011 the focus of production will be on the reprocessing of the final data run. After the reprocessing, the simulation production and the skimming will follow.

BaBar contributes about 50% of the SLAC general queues capacity. Data production is still done on *BaBar* dedicated machines. This will ramp down to 25% of its original allocation by the end of 2012. Support manpower is projected to flatten out at 0.5 FTE at that time.

Fermi

The Fermi telescope was launched in June 2008 and began sending back data 2 weeks later. It has been running smoothly since then, delivering some 13 GB of compressed raw science data daily. SLAC has responsibility for processing the data into high-level analysis products though transfer to the NASA Science Support Center. The effort also includes machinery to perform data quality and trending. SLAC also acts as the distribution center for the data and software to the collaboration. Prior to launch, the collaboration engaged in several Data Challenges with, first, a detailed sky model for the science groups to unravel, and then full scale operations simulations to practice actual data receipt and turn-around. Once real data arrived, it all just worked. Additionally SLAC has played a key role in developing and maintaining the simulation, reconstruction and high level analysis software.

The data handling tools were built with an emphasis on web access – the success of which is illustrated by the fact that all data processing and monitoring is done over the web. The key tools built were a workflow engine, with which we manage tens of thousands of batch jobs per day in parallel tasks; a data catalogue to organize and fetch registered data; data quality, monitoring and trending tools; smart data servers to deliver the astrophysics data, e.g., by location on the sky; code distribution gui's. Other collaborative tools were personnel database interfaces and management of Fermi's process for tracking and approving presentations.

EXO

The Enriched Xenon Observatory is an experiment to search for "neutrino-less double beta decay" using large amounts of xenon isotropically enriched in the isotope 136. A 200-kg detector (EXO200) using liquid Xe is now installed at the Waste Isolation Pilot Plant (WIPP) near Carlsbad, New Mexico.

For the EXO200 experiment SLAC is responsible for data acquisition and on-line monitoring as well as data transportation from WIPP to SLAC. Due to the lack of high-speed network connections to the experiment site data is transported to SLAC using removable 1.5TB disks. Reconstruction of all incoming data is performed at SLAC reusing the automated pipeline and data catalog system developed for the Fermi Gamma Ray Space Telescope. SLAC is also the hub for collaboration data analysis and data archiving.

A number of web tools have been developed for the EXO200 experiment, in many cases by customizing similar tools developed at SLAC for Fermi. Tools include a collaboration database, online data quality monitoring and trending system, online event display and data acquisition control system. EXO200 also makes extensive use of other web based collaborative tools hosted at SLAC, including the confluence wiki system.

A series of engineering runs were taken at the end of December with the TPC full of natural Xe. All production systems were employed to provide the first full-scale test of the system, producing about 25TB of data in two weeks. This data was successfully transported to SLAC and processed using the automated pipeline, and first analyses of this calibration data are now underway.

ATLAS

SLAC PPA has overall responsibility for the ATLAS Tier 2 center at SLAC (which is operated by the Computing Division) whereas the PPA/SCA department has direct responsibility for the other ATLAS M&O-funded personnel supporting ATLAS computing.

The Tier 2 is a competitive member of the US ATLAS cloud in terms of availability and the quantity of delivered CPU and storage. The Tier 2 funding and personnel also provide expert xrootd-focused support to US and International ATLAS. The cost of this support is covered by about 20% of the Tier 2 budget.

In addition to the closely connected Tier 2 and xrootd responsibilities, the PPA/SCA group has also been providing expert support for ATLAS for the Production System.

Accelerator Modeling through Advanced Computations

With support from SLAC's Accelerator Research Division and the SciDAC program, the Advanced Computations effort benefits the national and international accelerator community by developing and supporting advanced electromagnetic simulation capabilities for the RF design of existing and future facilities in the US and beyond.

Parallel code development. A suite of parallel electromagnetic codes, **ACE3P**, based on the higher-order finite element method, has been developed to provide solutions to challenging problems in accelerators. **ACE3P** has demonstrated the capability to simulate RF cavities and structures with an accuracy and at a scale previous not possible and covers a broad range of accelerator applications. The code suite is widely used by the accelerator community with a user base covering many accelerator Labs, universities and industry. **ACE3P** has the potential to replace commercial software thus lowering the project cost throughout the DOE accelerator complex.

Computational science R&D under SciDAC. R&D on computer science and applied math is an essential component of SLAC's electromagnetic effort under SciDAC. This activity facilitates the collaboration with the SciDAC Centers for Enabling Technologies & Institutes to develop the technologies that enable large-scale simulations. One such collaboration in experimental diagnosis, advanced computing and applied math research solved the beam breakup problem at CEBAF which has been widely recognized as a notable SciDAC success.

Accelerator modeling and simulation. **ACE3P** has been applied to a wide range of accelerator projects in Accelerator Science and Development. Recent applications include wakefield simulation for the CLIC Power Extraction & Transfer Structure (PETS) and accelerator structure, compact crab cavity and rotatable collimator designs for LHC, main injector cavity design for Project-X, dark current studies in high gradient structures, and power coupling in photonic bandgap structures for dielectric laser acceleration.

High performance computing for large-scale simulation. Leveraging the DOE resources at NERSC and NCCS, the accelerator modeling undertaken by Advanced Computations has made an impact on the ILC design by carrying out one of the largest wakefield and trapped mode calculations for a cryomodule consisting of 8 superconducting cavities. The simulation used 1024 cores/300 GBytes on NERSC Seaborg in the frequency domain, and 4096 cores on NCCS Jaguar in the time domain. Maintaining **ACE3P's** capabilities on evolving computing platforms remains a challenge.

KIPAC Computing

KIPAC, a joint SLAC/Stanford Institute, hosts a wide array of science programs. The SLAC based KIPAC computing department develops and maintains the computing environment used for analysis, simulation, storage, archiving, visualization and data-access. This work is carried out with assistance and support from the SLAC Computing Division.

The KIPAC research community includes ~150 students, postdocs, faculty and staff spanning ~10 groups and projects. KIPAC computing has supported LSST, DES, CDMS, Fermi, along with faculty-based research in cosmology. KIPAC computing maintains web based documentation and collaboration services. KIPAC staff maintain ~100 open source software packages useful for research computing. In addition to generic SLAC computing resources, KIPAC maintains several specialized systems.

KIPAC has some specialized computing needs: SMP work and a 3D projection system, used primarily to visualize the time evolution of cosmological simulations. KIPAC is developing additional data access and visualization capabilities with a tiled display wall and 10 Gbps networking between data sources and visualization resources.

Both n-body and hydrodynamic codes are heavily used by the computational cosmologists at KIPAC. These programs support the investigations into dark energy, dark matter, and the connection of fundamental physics to cosmology. The work on galaxy clusters as cosmology probes also rely on these systems, as well as realistic mock galaxy catalogs for DES. The hydro code ENZO has been significantly extended at SLAC to include new physics and improved dynamic range. Work is ongoing to enable scaling to much larger clusters. Also, these systems have been tuned to give highest priority to the larger MPI style jobs, with the result that embarrassingly parallel jobs supporting LSST, CDMS, and others can run unconstrained during periods of low usage. The visualization systems are a developing tool supporting analysis of simulations as well as education and public outreach.

Scientific Computing Applications Department

A merger of PPA's Fermi Data Handling team with the GEANT4 and xrootd teams was undertaken in spring 2010 to form the Scientific Computing Applications (SCA) Department in PPA. The goals were to leverage tools developed for Fermi and the Linear Collider detector studies to support the several small scale experiments that now characterize PPA: EXO, SCDMS, HPS, and Fermi.

SCA Data Handling Group

The tools are also applicable to larger experiments (e.g., LSST) and potentially outside PPA (e.g., LCLS). SCA also provides advice to the PPA Director in areas of computing and promotes communication and co-operation amongst the various PPA projects. It also helps manage the migration of personnel onto new projects.

The SCA data handling group has been formed to develop and deploy data handling, visualization and analysis tools for use by different experiments at SLAC, to encourage the efficient reuse of tools between different experiments at SLAC, and to invest in new computing technologies for use on future experiments. Where appropriate, the data handling group is building collaborations with other departments at SLAC and at other laboratories. Tools developed by the data handling group are in widespread use, although outreach beyond SLAC collaborations is severely limited by manpower.

GEANT4

GEANT4 is an object-oriented toolkit for the simulation of the interaction of particles with matter in complex geometries. GEANT4 is used by all LHC experiments, and by all SLAC HEP experiments. BaBar was the world's first major experiment to rely on GEANT4. This led SLAC set up its GEANT4 team that contributes to GEANT4 core development and supports SLAC programs.

The current SLAC GEANT4 team comprises the GEANT4 Spokesperson and former and original Chief Architect, the GEANT4 Hadronics Coordinator, the GEANT4 Visualization Coordinator, expertise in neutron and nuclear physics. The working model for the team is that members should spend 50% of their time on core GEANT4 development (funded by DOE/HEP/Computational HEP) and their remaining time on direct involvement in experiments' simulation activities funded by the experiments themselves. In addition to the management of the collaboration, the main focus of GEANT4 core development at SLAC has been on the needs of the LHC program:

- Precision and performance of hadronic models
- Architectural revision of GEANT4 kernel for performance and maintainability

Highlights of experiment support provided by the SLAC Geant4 team:

- ATLAS – muon system geometry; "stepper dispatcher" for efficient tracking in the ATLAS magnetic field; general code revisions for better performance; cavern background simulation (GEANT4 geometry + FLUKA)
- SuperCDMS – development of simulation system including simulation of phonons

- BaBar, FGST, ILC – ongoing support as needed.

The SLAC GEANT4 team has been engaged in outreach to HEP and non-HEP applications in the US and beyond. SLAC instituted the very popular idea of "Users Workshops and Tutorials", and has engaged with NASA supported work (Vanderbilt …), industry (Varian …), and medicine (NIH …). In general, the US non HEP applications do fund the work that benefits them directly, but the relationships have not yet reached the point of bringing in funding for core GEANT4.

xrootd

The xrootd system implements a structured clustering architecture to provide low latency access to data. In essence, xrootd clusters arbitrary data servers to form what appears to be a single server with a uniform name space but does so in a way that virtually has no limit and yet bounds the clustering overhead to $\log_{64}(n)$ time. The architecture also allows clusters to be federated across geographical boundaries to produce super-clusters. This provides opportunities for real-time access to vast amounts of data with minimal overhead; either machine or human resources.

Access to large quantities of data for geographically distributed groups is an innate High Energy Physics (HEP) problem as physics collaborations are dispersed and rather large. Because xrootd is synergistic with such an environment, it has become quite popular in HEP. With collaborative support between a team at CERN and SLAC, where it was originally developed, xrootd supports a wide array of data analysis today:

- Primary LHC data analysis at CERN,
- Global data analysis across all ALICE clusters,
- US ATLAS and CMS Tier 3 analysis site federation,
- Large scale opportunistic data analysis clustering at RHIC/BNL, and
- Local data analysis at many sites to support the Antares, ATLAS, BaBar, CMS, Compass, dchooz, EXO, Fermi, Hess, Indra, Opera, Panda, and Virgo experiments.

xrootd is also being used as an embedded system in the Parallel Root Facility (PROOF) for parallel data analysis and the Disk Pool Manager (DPM) to provide efficient data access. Both systems are widely distributed across the HEP community. Indeed, xrootd is available via the OSG/VDT and CERN Root/PROOF distribution channels.

LCSim

As the complexity and resolution of modern detectors increases, the need for detailed simulation of the experimental setup also becomes more critical. Designing the detectors requires efficient tools to simulate the detector response and reconstruct the events. We have developed efficient and flexible tools for detailed physics and detector response simulation as well as event reconstruction and analysis. The primary goal has been to develop a software toolkit and computing infrastructure to allow physicists from universities and labs to quickly and easily conduct physics analyses and contribute to detector research and development.

The lcsim project is composed of two distinct, but related applications. We have developed the detector response package, slic, which harnesses the full power of the GEANT4 toolkit without requiring the end user to have any experience with either GEANT4 or C++, thereby allowing the user to concentrate on the physics of the detector system. The system is powerful, yet simple to install and flexible enough to accommodate new detector geometries and technologies. All of the detector properties, not just the geometry, are definable at runtime with an easy to use xml format.

The output simulated detector response is provided in a simple, well-documented format (LCIO) allowing further processing or event reconstruction to be undertaken with a minimum of effort. LCIO is both an event data model and persistency framework, which provides base classes appropriate for most HEP applications. It was developed and is maintained as a joint SLAC/DESY venture, with contributions from other institutions. It is performant, with on-the-fly data compression and random access, well documented, and has C++, Java, python and FORTRAN bindings.

The Java based toolkit org.lcsim was developed for full event reconstruction and analysis. Java provides both a very powerful object-oriented language for development and transparent cross-platform "write once, run anywhere" portability. The components are fully modular and are available for tasks from digitization of tracking detector signals through to cluster finding, pattern recognition, track fitting, and analysis. The combination of an easy to use and flexible framework and the user-friendliness of Java has enabled contributions from a large number (~100) of both developers and users over a broad range of functionality.

The main focus of the lcsim core development team at SLAC has been on serving the needs of the ILC program, and the Silicon Detector (SiD) in particular. It was used to demonstrate the physics capabilities of SiD, resulting in the successful validation of the SiD concept during the Letter Of Intent exercise. Other users include:

- CERN CLIC physics and detector community
- FNAL Muon collider physics and detector community
- JLAB heavy-photon search proposals
    - HPS: fixed target, forward detector
    - Darklight: gas-jet, asymmetric 4π detector
- FNAL dual-readout crystal calorimetry

The SLAC lcsim team has been engaged in outreach to HEP applications in the US and beyond through software workshops, on-site tutorials, and extensive online documentation. Outreach to other application domains has taken place in non-HEP-specific conferences and workshops (e.g. IEEE NSS, IWORID, etc.)